Instrumenting an SMT Solver to Solve Hybrid Network Reachability Problems

Daniel Bryce
SIFT, LLC.,
dbryce@sift.net

Sergiy Bogomolov
IST Austria,
sergiy.bogomlov@ist.ac.at

Alexander Heinz
University of Freiburg,
alexander.heinz@saturn.uni-freiburg.de

Christian Schilling
University of Freiburg,
schillic@informatik.uni-freiburg.de

arXiv:1609.03847v1 [cs.AI] 13 Sep 2016

***Abstract*—PDDL+ planning has its semantics rooted in hybrid automata (HA) and recent work has shown that it can be modeled as a network of HAs. Addressing the complexity of nonlinear PDDL+ planning as HAs requires both space and time efficient reasoning. Unfortunately, existing solvers either do not address nonlinear dynamics or do not natively support networks of automata.**

We present a new algorithm, called `HNSolve`, which guides the variable selection of the dReal Satisfiability Modulo Theories (SMT) solver while reasoning about network encodings of nonlinear PDDL+ planning as HAs. `HNSolve` tightly integrates with dReal by solving a discrete abstraction of the HA network. `HNSolve` finds composite runs on the HA network that ignore continuous variables, but respect mode jumps and synchronization labels. `HNSolve` admissibly detects dead-ends in the discrete abstraction, and posts conflict clauses that prune the SMT solver's search. We evaluate the benefits of our `HNSolve` algorithm on PDDL+ benchmark problems and demonstrate its performance with respect to prior work.

I. Introduction

Recent planners ([7] and [6]) for PDDL+ [19] represent actions, processes, events, and state variables as a network of synchronized hybrid automata (HA), but there are no suitable algorithms for reasoning about nonlinear change in a network of automata. We address nonlinear PDDL+ problems by adapting the dReal Satisfiability Modulo Theories (SMT) solver [10], [22], which has been previously shown to address nonlinear PDDL+ as a single hybrid automaton. It is well-known that reasoning about the explicit parallel composition of a network of automata as a single automaton is usually a poor choice because it grows exponentially in the size of the individual synchronized automata [2], [11], [14]. We base our HA network SMT encoding upon prior work [14] that represents each automaton independently and adds synchronization constraints. We show that a direct network encoding results in better scalability in the number of encoding steps than a single automaton encoding if it is also matched with an appropriate search strategy.

We extend and reinterpret the approach taken by Bryce *et al.* [10] within our proposed `HNSolve` algorithm. They guide dReal to systematically select variable assignments that correspond to discrete-feasible runs of a hybrid automaton (i.e., ignoring continuous variables). In this manner, they force dReal to perform a heuristic depth-first search that considers all runs through the hybrid automaton. This led to increased scalability because dReal did not consider variable assignments that correspond to discrete-infeasible runs. A similar technique is employed by the dReach [22] and BACH [12] algorithms, but these encode a different SMT instance for each discrete run and do not benefit from learned conflict clauses.

`HNSolve`, like the work of Bryce *et al.* [10], guides dReal variable selection to construct discrete-feasible runs, but differs in two main aspects. First, `HNSolve` constructs composite runs for a network of HAs and not a single HA. Each composite run represents a feasible sequence of synchronized transitions within each automaton. `HNSolve` also performs a heuristic depth-first search over the possible composite runs. The second difference is that `HNSolve` also learns conflict clauses that it adds to the SMT encoding. While searching the space of discrete-feasible composite runs, `HNSolve` may be unable to find a run that extends the current partial run. dReal's current variable assignment encodes a partial run prefix, and the conflict clause blames the current variable assignment encoding the prefix. We find that adding conflict clauses is a critical new aspect of our approach, and demonstrate that it improves solver scalability as the number of steps in the encodings increase. Developing problem-specific SMT solver algorithms was also recently explored in the context of program analysis [25]. We hope to advance the understanding of how problem-specific, and arguably invasive, modifications to SMT solvers can bring about improved performance.

We evaluate the `HNSolve` algorithm on several PDDL+ planning benchmarks encoded as networks of hybrid automata, as developed in prior work [7]. We compare flat, precompiled encodings of the networks to our direct encoding of the networks. We show that our encoding outperforms encodings based upon explicit parallel composition and is competitive with the state of the art in PDDL+. We also demonstrate the advantages of using the `HNSolve` algorithm to improve performance.

II. Hybrid Automata Background

We discuss how to represent HA networks in the $\mathcal{L}_{\mathbb{R}_\mathcal{F}}$ language and their semantics, as follows.

First-order Theories of the Reals: The $\mathcal{L}_{\mathbb{R}_\mathcal{F}}$ language represents the first-order signature over the reals with the set $\mathcal{F}$ of computable real functions.

Definition 1 ($\mathcal{L}_{\mathbb{R}_\mathcal{F}}$-Formulas)**.** *$\mathcal{L}_{\mathbb{R}_\mathcal{F}}$-formulas are first-order formulas over real numbers, whose signature allows an arbi-*

trary collection $\mathcal{F}$ of Type 2 computable real functions [21]. The syntax is:

$$t := c \mid x \mid f(t(\vec{x}));$$
$$\varphi := t(\vec{x}) > 0 \mid t(\vec{x}) \geq 0 \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \exists x_i \varphi \mid \forall x_i \varphi.$$

A function is Type 2 computable if it can be algorithmically evaluated up to an arbitrary numerical accuracy. All common continuous real functions are Type 2 computable.

Networks of Hybrid Automata: We solve HA network reachability problems by expressing the network in $\mathcal{L}_{\mathbb{R}_\mathcal{F}}$ and then unrolling it over k steps and associating a goal region $\mathsf{goal}(\vec{x}_k^t)$. The goal is also an $\mathcal{L}_{\mathbb{R}_\mathcal{F}}$ formula of the form $\mathsf{mode}_k \wedge \varphi(\vec{x}_k^t)$ where mode_k defines the HA modes to reach and $\varphi(\vec{x}_k^t)$ defines the variable values to reach.

Definition 2 ($\mathcal{L}_{\mathbb{R}_\mathcal{F}}$-Representations of Hybrid Automata)**.** *A hybrid automaton in $\mathcal{L}_{\mathbb{R}_\mathcal{F}}$-representation is a tuple*

$$H = \langle X, Q, \mathsf{flow}, \mathsf{inv}, \mathsf{jump}, \mathsf{init}, \mathsf{L} \rangle$$

where

- $X \subseteq \mathbb{R}^n$ *for some* $n \in \mathbb{N}$*,*
- $Q = \{q_1, \ldots, q_m\}$ *is a finite set of modes,*
- $\mathsf{flow} = \{\mathsf{flow}_q(\vec{x}^0, \vec{x}^t, t) : q \in Q\}$ *is the set of ODEs describing the flow for each mode,*
- $\mathsf{inv} = \{\mathsf{inv}_q(\vec{x}) : q \in Q\}$ *is a set of invariants for each mode,*
- $\mathsf{jump} = \{\mathsf{jump}_{q \xrightarrow{\ell} q'}(\vec{x}^t, \vec{x}^{0'}) : q, q' \in Q\}$*, where each element is a transition from mode q to q' using the synchronization label set ℓ. Each formula $\mathsf{jump}_{q \xrightarrow{\ell} q'}(\vec{x}^t, \vec{x}^{0'})$ is of the form $\varphi(\vec{x}^t) \wedge \psi(\vec{x}^t, \vec{x}^{0'})$, where $\varphi(\vec{x}^t)$ is a conjunction specifying the guard, $\psi(\vec{x}^t, \vec{x}^{0'})$ is a conjunction specifying the discrete update,*
- $\mathsf{init} = \{\mathsf{init}_q(\vec{x}) : q \in Q\}$ *is the set of initial states, and*
- L *is a finite set of synchronized event labels.*

Gao *et al.* [22] describe how to unroll this encoding for a single automaton. The important aspects of the unrolling are to time stamp each of the continuous variables to denote their value at the start ($\vec{x}_i^0$) and end ($\vec{x}_i^t$) of the mode at step i, as well as the time t_i spent in step i.

Definition 3. *(HA Runs) Each run τ on H is a series of states*

$$(t_0, q_0, s(\vec{x}_0^0), s(\vec{x}_0^t)),$$
$$(t_1, q_1, s(\vec{x}_1^0), s(\vec{x}_1^t)),$$
$$\ldots,$$
$$(t_k, q_k, s(\vec{x}_k^0), s(\vec{x}_k^t))$$

where t_i is the time spent in step i, q_i is a mode, and $s(\vec{x}_i^0)$ and $s(\vec{x}_i^{t_i})$ are respective valuations on $\vec{x}_i^0$ and $\vec{x}_i^t$ upon entering or leaving q_i.

The init set defines what constitutes a legal initial mode q_0 and state valuation $s(\vec{x}_0^0)$. The flow and inv sets define legal pairs of state valuations $s(\vec{x}_i^0)$ and $s(\vec{x}_i^t)$ and occupancy times t_i. The jump set defines legal pairs of modes q_i and q_{i+1}, state valuations $s(\vec{x}_i^t)$ and $s(\vec{x}_{i+1}^0)$, and occupancy times t_i.

Definition 4. *(Hybrid Automaton Network) A network $\mathcal{N} = \{H_1, \ldots, H_n\}$ of hybrid automata is a set of hybrid automata.*

Definition 5. *(HA Network Runs) Each run τ on $\mathcal{N}$ is a series of composite states*

$$(t_0, \vec{q}_0, s(\vec{x}_0^0), s(\vec{x}_0^t)),$$
$$L_0,$$
$$(t_1, \vec{q}_1, s(\vec{x}_1^0), s(\vec{x}_1^t)),$$
$$L_1,$$
$$\ldots,$$
$$L_{k-1},$$
$$(t_k, \vec{q}_k, s(\vec{x}_k^0), s(\vec{x}_k^t))$$

interleaved with sets of synchronization labels L_i, where each state includes a vector of modes $\vec{q}_i$ and each $\vec{x}_i^0$ and $\vec{x}_i^t$ is a valuation on $X = X_1 \cup \ldots \cup X_n$.

In order for a HA network run to be legal, it must be consistent with each of the init, flow, and inv sets of each individual HA. Each label set L_i determines the legal constituent jumps that make up a composite jump. In a composite jump, each automaton either changes modes as defined by its jump set, or remains in the same mode. Let $\mathsf{sync} = \ell \cap \mathsf{L}$ be the set of labels a jump must synchronize upon. If the composite label is L_i, the automaton must take one of the jumps from the current mode where $\mathsf{sync} \subseteq L_i$. If no jump may be taken and $\mathsf{L} \not\subseteq L_i$, then the automaton may remain in the same mode.

III. Hybrid Network Solver

We describe an extension called the hybrid network solver (`HNSolve`) to dReal's SMT framework that customizes the problem solver for reachability checking in networks of hybrid automata.

dReal: dReal checks whether an $\mathcal{L}_{\mathbb{R}_\mathcal{F}}$ formula is δ-satisfiable (a decidable problem) by combining a SAT solver [18] with an ICP solver.[1] dReal employs the DPLL(T) framework [8] for SMT. It first solves the Boolean constraints to find a satisfying set of literals of the form $(t(\vec{x}) \geq 0)$ or $\neg(t(\vec{x}) \geq 0)$. This conjunctive set of literals imposes a set of numeric constraints that are solved using ICP. If successful, dReal finishes, and otherwise, the ICP solver returns a set of literals that explain inconsistency. The inconsistent literals become a conflict clause that can be used by the SAT solver. If the SAT solver cannot find a satisfying set of literals, then it returns with an unsatisfiability result.

The ICP solver uses the branch and prune [27] algorithm to refine a set of intervals over the continuous variables (called a box). Each branch splits the interval of a single continuous variable, creating two boxes. Pruning operators propagate the constraints to shrink the boxes. ICP continues to branch and prune boxes until it finds a box that is δ-satisfiable or establishes that no such box exists (i.e., the constraints are inconsistent). A box is δ-satisfiable when for any vector of

[1] www.ibex-lib.org

values $\vec{x}$ represented by the box, each constraint $f(\vec{x}) \geq -\delta$ is satisfied.

Hybrid Network Solver: `HNSolve` wraps the SAT solver and ICP solver by suggesting variable assignments and adding conflict clauses. Where dReal would normally make decisions without realizing it is solving a hybrid network reachability problem, `HNSolve` organizes its search around potentially viable runs on the network. To see this, consider a situation where dReal has made several variable assignments and decided that $\vec{q}_0$ and $\vec{q}_k$ are the initial and final modes of a run satisfying the $\mathsf{goal}(x_k^t)$. For the sake of example, the first automaton H_1 does not include jumps that connect q_0 to q_k. A simple shortest path algorithm can recognize that there is no path from q_0 to q_k, but dReal will continue to make assignments unnecessarily and eventually backtrack.

`HNSolve` solves a discrete abstraction [1], [11] of the hybrid reachability problem and coordinates the SAT solver in making assignments corresponding to its solution. Its approach is to find a sequence of mode vectors and synchronization labels of the form

$$\vec{q}_0, L_0, \vec{q}_1, L_1, \ldots, L_{k-1}, \vec{q}_k$$

that are consistent with the jump transitions and synchronize. `HNSolve` then extracts the literal assignments corresponding to this sequence and guides the SAT solver to realize this run. If the SAT or ICP solvers discover an inconsistency, then `HNSolve` rebuilds the run from the point to which the SAT solver backtracks.

`HNSolve` interfaces with the SAT solver through three main methods: $\mathsf{getTrail}()$, $\mathsf{assertLit}()$ and $\mathsf{assertClause}()$. The $\mathsf{getTrail}()$ method returns the SAT solver's assignment stack, including all premises, decision variable assignments, and inferred variable assignments. The $\mathsf{assertLit}()$ and $\mathsf{assertClause}()$ methods assert a literal or clause assignment, respectively, and return $unsat$, δ-sat, $consistent$, or *backtrack*. We note that there is no need for a corresponding $\mathsf{retractLit}()$ method because backtracking is handled by the SAT solver as part of the calls to $\mathsf{assertLit}()$ and $\mathsf{assertClause}()$.

`HNSolve`, Algorithm 1, first computes the cost (described below) to reach each mode in each automaton in line 2 and then encodes the problem in $\mathcal{L}_{\mathbb{R}_\mathcal{F}}$ in line 3. The algorithm works in two phases. In lines 5 - 14, the solver generates a run suffix $\mathcal{P}$. Each run is represented by a set of literals of the form $\{\mathsf{mode}_0^1 = q_0^1, \mathsf{mode}_0^2 = q_0^2, \ldots, \mathsf{sync}_0^l, \ldots, \mathsf{mode}_k^1 = q_k^1, \mathsf{mode}_k^2 = q_k^2, \ldots\}$, defining the modes of each automaton at each step and the synchronization labels of each transition. If the current run represented by the SAT solver's trail cannot be extended, then $\mathsf{genRun}()$ returns a "nil" run. Lines 8 - 12 create a conflict clause from the decision variables on the SAT solver's trail (i.e., negate the corresponding literals) and assert the new clause. The $\mathsf{assertClause}()$ method returns $unsat$ when the SAT solver determines the new clause causes unsatisfiability; otherwise, the SAT solver will backtrack as part of $\mathsf{assertClause}()$ and we will find a new run on line 6.

Upon finding a non-nil run, lines 15 - 26 assert each literal in the run from the beginning of the run to the end (i.e., following a run forward to one of the goal modes of each automaton). After successfully asserting each literal on the run, `HNSolve` asserts the nil literal, which signals the SAT solver to complete any remaining assignments. If the run leads to a δ-sat solution, then the algorithm returns; otherwise, if asserting a literal is not consistent (i.e., returns *backtrack* or $unsat$), then the solver attempts to find a new run on line 6.

The $\mathsf{genRun}()$ method (Algorithm 2) finds a run on the network that is consistent with the current SAT solver assignment T. It uses depth-first search (line 3 and Algorithm 3) to find a search stack S. The search stack S includes a transition for each automaton for each step 0 to k-1. From S, the genRun method extracts the literals needed to encode the run (lines 7-14). The literals include mode choices (line 10) and synchronization labels (lines 11-13).

The depth-first search (Algorithm 3) generates the search stack S corresponding to a run on the network. It selects an initial mode of H_1 at depth 0, an initial mode of H_2 at depth 1, and so on. It selects a jump from the initial mode of H_1 (chosen at depth 0) at depth $|\mathcal{N}|$, a jump for H_2 at depth $|\mathcal{N}| + 1$, and so on. Thus, the first $\mathcal{N}$ levels of the stack correspond to initial modes, the second $|\mathcal{N}|$ levels to the zeroth step, and similarly for later steps. Lines 8-18 generate the

Algorithm 1: HNSolve algorithm.

```
1  HNSolve (N, G, k, M)
   input : A network of automata N, a reachability
           property G, and step and delay bound k and M.
2  cost ← getRunCosts(N);
3  encodeSMT(N, G, k, M);
4  while true do
5      repeat
6          P ← genRun(N, getTrail(), cost, k);
7          if P = nil then
8              C ← ⋁_{l∈decisions(getTrail())} ¬l;
9              res ← assertClause(C);
10             if res = unsat then
11                 return unsat
12             end
13         end
14     until |P| ≠ nil;
15     for l from 0 upto |P| do
16         if l < |P| then
17             res ← assertLit(P[l]);
18         else
19             res ← assertLit(nil);
20         end
21         if res = δ-sat then
22             return δ-sat;
23         else if res ≠ consistent then
24             break;
25         end
26     end
27 end
```

Algorithm 2: genRun algorithm.

```
1  genRun(𝒩, T, cost, k)
   input : A network of automata 𝒩, a stack of literals T,
           a cost function cost, and the step bound k.
2  𝒫 ← [];
3  S ← dfs(𝒩, [], cost, k);
4  if S = fail then
5  |  return nil;
6  else
7  |  for j = 0 ... |S| − 1 do
8  |  |  (q_i −ℓ_i→ q_i') ← S.get(j);
9  |  |  step ← j/|𝒩|;
10 |  |  𝒫.append((mode^i_step = q_i', ⊤));
11 |  |  for l ∈ ℓ do
12 |  |  |  𝒫.append((sync^l_step, ⊤));
13 |  |  end
14 |  end
15 end
16 return 𝒫;
```

Algorithm 3: Depth-First Search Algorithm.

```
1  dfs(𝒩, S, cost, k)
   input : A network of automata 𝒩, a search stack S, a
           mode cost function cost, and a step bound k.
2  if |S| = |𝒩|(k + 1) then
3  |  return S;
4  end
5  step ← |S|/|𝒩|;
6  i ← |S|%|𝒩|;
7  succ ← {};
8  if step = 0 then
9  |  for init_{q'}(x⃗) ∈ init_i do
10 |  |  succ ← succ ∪ (nil −{}→ q');
11 |  end
12 else
13 |  (q'' −ℓ'→ q) ← S.get(|S| − |𝒩|);
14 |  for jump_{q −L→ q'}(x⃗^t, x⃗^0) ∈ jump_i do
15 |  |  succ ← succ ∪ (q −ℓ→ q');
16 |  end
17 |  succ ← succ ∪ (q −{}→ q);
18 end
19 succ ←sort(filter(succ, i, 𝒩, S, step), cost);
20 for q −ℓ→ q' ∈ succ do
21 |  S.push(q −ℓ→ q');
22 |  if dfs(𝒩, S, k) ≠ fail then
23 |  |  return S;
24 |  else
25 |  |  S.pop();
26 |  end
27 end
28 return fail;
```

successors, line 19 filters the successors (described below), and lines 20-27 conduct the recursive step of the search.

Algorithm 4 removes successors that are either inconsistent with the SAT solver's current assignment (lines 8-9) or do not synchronize with the "sibling" jumps previously chosen for the current step (lines 13-20). The synchronization check involves three cases. The first and second case are for when either the current automaton or the sibling will persist its mode (i.e., not synchronize). We require that for a mode q_i to persist in automaton $\mathcal{H}_i$, the automaton must not be compelled to synchronize with its siblings at the current time step. More formally, the mode persistence is allowed if the public labels $\mathsf{L}_j \cap \ell_j$ for each chosen sibling jump $q_j \xrightarrow{\ell_j} q'_j$ do not intersect with the labels L_i. The third case checks that two jumps agree on publicly communicated labels. The jumps in $succ'$ are those that are possible in the discrete sense, but may not be possible if we were to consider their continuous variables in the guards, updates, or mode invariants.

After filtering the possible successors, the depth-first search will sort the jumps by increasing cost, where cost is defined by the successor mode's value in $cost$. The cost of each mode is defined by the minimum number of steps from an initial mode:

$$cost(q) = \begin{cases} 0 & : \mathsf{init}_q(\vec{x}) \in \mathsf{init}_i \\ \min\limits_{\mathsf{jump}_{q' \xrightarrow{\ell} q} \in \mathsf{jump}_i} cost(q') + 1 & : \text{otherwise} \end{cases}$$

In the next section, we detail the $\mathcal{L}_{\mathbb{R}_\mathcal{F}}$ encoding that we use to express the network HA reachability problem. We follow with a section describing how we experiment with `HNSolve` by omitting it entirely (i.e., use dReal alone), omitting lines 6-12 of Algorithm 1 to avoid learning conflict clauses, or using it in its entirety. We note that omitting conflict clauses results in an algorithm similar to that described by Bryce *et al.* [10], aside from our generalization to a network of automata.

IV. Network Encoding

We encode the parallel composition of a network of hybrid automata implicitly, as follows. We encode the mode at step i of each automaton with literals of the form $\mathsf{mode}_i^1 = q_1, \ldots, \mathsf{mode}_i^m = q_m$. We constrain the possible composite jumps with synchronized jump constraints, and noops. With these constraints, we avoid pre-computing all possible $O(2^{|\mathsf{jump}|})$ parallel jumps per step. Instead, we encode $O(|\mathsf{jump}|m)$ synchronized jump constraints and $O(|Q|m)$ noops. To determine which jumps must synchronize, we introduce literals for each label and constrain their values with the appropriate jumps. Noop (stutter) clauses encode cases where an automaton does not synchronize any of its transitions, and its mode persists.

We define the implicit parallel composition for a k-step M-delay reachability problem as the conjunction of clauses describing each individual automaton, and the goal:

Algorithm 4: Filter algorithm.

filter($succ$, i, $\mathcal{N}$, S, $step$)
input : A set of transitions $succ$, an index of the current automaton i, a network of automata $\mathcal{N}$, a search stack S, and the current time step $step$.
$succ' \leftarrow \{\}$;
$siblings \leftarrow \{\}$;
for $j = 0 \dots |S|\%|\mathcal{N}|$ **do**
  $siblings \leftarrow siblings \cup S.\text{get}(|\mathcal{S}| - (i - j))$;
end
for $(q_i \xrightarrow{\ell_i} q_i') \in succ$ **do**
  if $T.contains(mode^i_{step} = q_i',\ \bot)$ **or**
    $\exists l \in \ell_i.\ T.contains(sync^l_{step},\ \bot)$ **then**
    continue;
  end
  $syncs \leftarrow \top$;
  for $(q_j \xrightarrow{\ell_j} q_j') \in siblings$ **do**
    if $(\ell_j = \{\}$ **and** $q_j = q_j'$ **and** $\mathsf{L}_i \cap \mathsf{L}_j \cap \ell_i \neq \{\})$ **or**
      $(\ell_i = \{\}$ **and** $q_i = q_i'$ **and** $\mathsf{L}_i \cap \mathsf{L}_j \cap \ell_j \neq \{\})$ **or**
      $(\mathsf{L}_i \cap \mathsf{L}_j \cap \ell_i \neq \mathsf{L}_i \cap \mathsf{L}_j \cap \ell_j)$ **then**
      $syncs \leftarrow \bot$;
      break;
    end
  end
  if $syncs$ **then**
    $succ' \leftarrow succ' \cup (q_i \xrightarrow{\ell_i} q_i')$
  end
end
return $succ'$;

$$\exists^X \vec{x}_0^0\, \exists^X \vec{x}_0^t \dots \exists^X \vec{x}_k^0\, \exists^X \vec{x}_k^t$$
$$\exists^{[0,M]} t_0 \dots \exists^{[0,M]} t_k.$$
$$\left(\bigwedge_{j=1}^{n} \mathsf{autom}(H_j, k)\right) \wedge \mathsf{goal}(\vec{x}_k^t) \wedge \left(\bigwedge_{i=0}^{k-1} \bigvee_{l \in \mathsf{L}_1 \cup \dots \cup \mathsf{L}_n} \mathsf{sync}_i^l\right)$$

The $\vec{x}_i^0$ and $\vec{x}_i^t$ variables denote the values of continuous variables at the start and end of step i. The t_i variables denote the duration of step i. The encoding ensures that each automaton behaves appropriately, the goal is satisfied, and at least one non-noop transition occurs in each step.

The $\mathsf{autom}(H_j, k)$ clauses define the behavior of each automaton H_j as:

$$\mathsf{init}_j(\vec{x}_0^0) \wedge \bigwedge_{i=0}^{k} \mathsf{maintain}_j(i) \wedge$$
$$\left[\bigwedge_{i=0}^{k-1} \left(\bigvee_{q \in Q_j} \mathsf{noop}(q, i)\right) \vee \bigvee_{\mathsf{jump}_{q \xrightarrow{\ell} q'}(\vec{x}_i^t, \vec{x}_{i+1}^0) \in \mathsf{jump}_j} \mathsf{trans}_j(\mathsf{jump}_{q \xrightarrow{\ell} q'}(\vec{x}_i^t, \vec{x}_{i+1}^0), i)\right]$$

which constrain the initial state, the continuous change in each mode at each step, and the transitions between steps.

The clause $\mathsf{init}_j(\vec{x}_0^0)$ constrains the initial values of the variables and the initial mode. It defines:

$$\bigvee_{\mathsf{init}_q(\vec{x}) \in \mathsf{init}_j} \mathsf{init}_q(\vec{x}_0^0) \wedge (\mathsf{mode}_0^j = q)$$

to constrain the assignments to $\vec{x}_0^0$ and the initial mode.

The $\mathsf{maintain}_j(i)$ clause defines how the flows and invariants of the automaton H_j govern continuous change:

$$\mathsf{flow}_j(\vec{x}_i^0, \vec{x}_i^t, t_i) \wedge$$
$$\forall^{[0,t_i]} t\ \forall^X \vec{x}_i\ (\mathsf{flow}_j(\vec{x}_i^0, \vec{x}_i, t) \rightarrow \mathsf{inv}_j(\vec{x}_i))$$

where we note that the nested universal quantifiers ensure that the invariant holds for the entire time the mode is occupied. The nested quantifiers are a unique aspect of our encoding that enables us to reason about nonlinear change [23].

The $\mathsf{flow}_j(\vec{x}_i^0, \vec{x}_i^t, t_i)$ clause defines

$$\bigwedge_{q \in Q_j} (\mathsf{mode}_i^j = q) \rightarrow \mathsf{flow}_q(\vec{x}_i^0, \vec{x}_i^t, t_i)$$

The $\mathsf{inv}_j(\vec{x}_i)$ clause enables the invariants of the active modes by defining:

$$\bigwedge_{q \in Q_j} (\mathsf{mode}_i^j = q) \rightarrow \mathsf{inv}_q(\vec{x}_i)$$

Noop clauses $\mathsf{noop}(q, i)$ model asynchronous behavior where the automaton does not synchronize, and define:

$$\left(\bigwedge_{l \in \mathsf{L}_j} \neg \mathsf{sync}_i^l\right) \wedge (\mathsf{mode}_i^j = q) \wedge (\mathsf{mode}_{i+1}^j = q)$$

Jump transition clauses $\mathsf{trans}_j(\mathsf{jump}_{q \xrightarrow{\ell} q'}(\vec{x}_i^t, \vec{x}_{i+1}^0), i)$ define how each jump must synchronize and constrain the variables and modes:

$$\left(\bigwedge_{l \in \ell \cap \mathsf{L}_j} \mathsf{sync}_i^l \wedge \bigwedge_{l \in \mathsf{L}_j \setminus \ell} \neg \mathsf{sync}_i^l\right) \wedge \mathsf{jump}_{q \xrightarrow{\ell} q'}(\vec{x}_i^t, \vec{x}_{i+1}^0) \wedge$$
$$(\mathsf{mode}_i^j = q) \wedge (\mathsf{mode}_{i+1}^j = q')$$

k	Dom	Inst	F	F+H	F+H+L	C	C+H	C+H+L	N	N+H	N+H+L
3 (3)	Gen	0	0.22	0.16	0.17	1.52	0.11	0.12	0.30	0.13	0.13
7 (5)	Gen	1	-	1.66	1.15	-	3.45	3.41	1.39	0.75	0.77
11 (7)	Gen	2	-	-	-	-	735.00	738.10	20.69	3.62	3.62
15 (9)	Gen	3	-	-	-	-	-	-	-	12.23	12.77
19 (9)	Gen	4	-	-	-	-	-	-	-	37.77	36.05
23 (13)	Gen	5	-	-	-	-	-	-	-	100.21	92.26
27 (17)	Gen	6	-	-	-	-	-	-	-	511.00	360.63
31 (19)	Gen	7	-	-	-	-	-	-	-	561.06	482.03
6	Car1	1	0.85	0.84	0.89	9.77	1.21	0.98	0.9	1.33	1.32
5	Car1	2	1.59	0.75	0.74	13.18	0.97	0.9	0.84	1.09	1.05
5	Car1	3	0.99	0.72	0.72	44.64	1.22	1.15	1.82	1.56	1.59
5	Car1	4	1.63	0.86	0.89	83.55	1.63	1.49	1.79	2.22	2.43
5	Car1	5	7.41	1.39	1.43	229.54	2.09	1.93	2.62	3.34	3.95
5	Car1	6	10.01	1.64	1.69	448.62	2.82	2.56	7.05	5.12	6.11
5	Car1	7	9.98	1.92	1.98	-	3.94	3.57	7.57	7.93	9.77
5	Car1	8	10.69	1.67	1.68	-	6.78	6.26	14.98	14.87	18.33
5	Car1	9	18.65	1.96	1.94	-	7.69	7.34	23.55	21.41	24.13
5	Car1	10	46.74	2.62	2.51	-	12.42	10.15	40.28	34.66	40.01
12	Car2	1	-	25.02	78.63	-	-	-	37.45	236.39	9.87
10	Car2	2	-	336.86	330.61	-	-	-	-	-	11.84
10	Car2	3	-		-	-	-	-	-	-	24.98
10	Car2	4	-		-	-	-	-	-	-	20.04
10	Car2	5	-		-	-	-	-	-	-	72.35
10	Car2	6	-		-	-	-	-	-	-	119.69
10	Car2	7	-		-	-	-	-	-	-	194.38
10	Car2	8	-		-	-	-	-	-	-	408.78
10	Car2	9	-		-	-	-	-	-	-	222.63
10	Car2	10	-		-	-	-	-	-	-	328.82

TABLE I
RUNTIME (S) ON LINEAR INSTANCES. "-" INDICATES A TIMEOUT.

V. EMPIRICAL EVALUATION

Our evaluation studies the effectiveness of our network encoding and `HNSolve` algorithm. Specifically, we compare three configurations of our solver on several hybrid automaton encodings of PDDL+ problems. The configurations include an unmodified dReach/dReal solver, the addition of the `HNSolve` algorithm without clause learning, and `HNSolve` with clause learning. We evaluate the configurations on single automaton encodings [10], and networks of automata encodings based on that of Bogomolov *et al.* [7]. With the network encodings, we either take the parallel composition and encode a single automaton, or encode the network as described in the previous section.

We note that the encodings used in prior work differ in whether they include a "lock" for the actions. Bryce *et al.* [10] hand-encode a single automaton for each problem that ensures no two actions can occur at the same time. In a network of automata, where each action is represented by its own automaton, Bogomolov *et al.* [7] ensure that no two actions occur, start, or end simultaneously by introducing a lock automaton. The network of automata models each action so that it must acquire and release the lock when it occurs (atomic actions) or starts/ends (durative actions). This causes the network encoding to require twice the number of encoding steps than the single automaton encoding used by Bryce *et al.* [10]. We notice that the two step lock is only necessary when enforcing ϵ-separation of the actions. Since Bryce *et al.* [10] do not model ϵ-separation we can match their required number of encoding steps with the network of automata by using a single lock transition that synchronizes with each action.

We also compare `HNSolve` with dReal and existing planners, including SpaceEx [3]–[5], [20], CoLin [15], and UPMurphi [17]. We reproduce previously published results [7] for the SpaceEx, CoLin, and UPMurphi approaches, but report runtimes for `HNSolve` and dReal from the same machine, a 2.6 GHz Intel Core i7 and 8GB RAM.

Our approach inherits some of the limitations of using SpaceEx with the Bogomolov *et al.* [7] network representation of PDDL+. The encoding does not respect the "must" semantics of PDDL+ wherein processes and events must occur when enabled. However, this limitation is not realized in our chosen benchmarks because any use of a process or event is advantageous to the plan. We also note that dReal (and `HNSolve` as a result) find δ-satisfiable solutions to the $\mathcal{L}_{\mathbb{R}_\mathcal{F}}$ encoding. Owing to the undecidable nature of nonlinear hybrid systems, dReal cannot guarantee that a δ-satisfiable solution, which bounds the values of the continuous variables, contains a realizable plan. Defining an appropriately small value for δ minimizes this concern. We also note that dReal in itself is not a full planner. We report results for the minimum step length required to find a plan. A number of strategies for exploring different step lengths in parallel or in sequence have been studied in SAT based planning and can be applied here. We note that these considerations must be incorporated when comparing the results for our approach with that of the other planners.

Domains: We use the Generator and Car domains from the literature [7] and the Dribble domain [10]. We compare on linear and nonlinear versions of Generator and Car, but only

k	Dom	Inst	F	F+H	F+H+L	C	C+H	C+H+L	N	N+H	N+H+L
3 (3)	Gen	0	0.15	0.16	0.14	1.40	0.14	0.15	0.49	0.49	0.18
7 (5)	Gen	1	0.76	1.66	1.77	-	4.21	4.22	1.74	1.16	1.14
11 (7)	Gen	2	26.36	-	-	-	879.36	876.92	173.78	6.08	5.78
15 (9)	Gen	3	-	-	-	-	-		310.70	22.91	21.77
19 (9)	Gen	4	-	-	-	-	-		-	69.16	66.89
23 (13)	Gen	5	-	-	-	-	-	-	-	320.21	316.45
27 (17)	Gen	6	-	-	-	-	-	-	-	530.77	1044.27
6	Car1	1	5.49	1.63	1.54	-	-	-	3.46	2.81	23.48
5	Car1	2	3.45	1.46	1.38	-	-	-	2.26	2.19	22.35
5	Car1	3	8.06	1.48	1.44	-	-	-	6.15	3.87	43.57
5	Car1	4	4.73	1.51	1.52	-	-	-	7.88	7.27	81.93
5	Car1	5	5.25	1.53	1.49	-	-	-	10.91	9.82	145.64
5	Car1	6	6.42	1.47	1.50	-	-	-	19.96	17.05	251.90
5	Car1	7	7.02	1.45	1.53	-	-	-	42.53	29.73	465.09
5	Car1	8	9.79	1.44	1.53	-	-	-	76.78	45.12	216.67
5	Car1	9	10.23	1.93	2.05	-	-	-	143.52	76.76	356.15
5	Car1	10	12.45	1.92	2.05	-	-	-	221.08	121.31	498.48
12	Car2	1	-	313.77	219.1	-	-	-	-	-	24.02
10	Car2	2	-	966.82	342.12	-	-	-	-	-	23.17
10	Car2	3	-	-	-	-	-	-	-	-	46.51
10	Car2	4	-	-	-	-	-	-	-	-	85.67
10	Car2	5	-	-	-	-	-	-	-	-	146.43
10	Car2	6	-	-	-	-	-	-	-	-	246.70
10	Car2	7	-	-	-	-	-	-	-	-	448.86
10	Car2	8	-	-	-	-	-	-	-	-	217.37
10	Car2	9	-	-	-	-	-	-	-	-	370.81
10	Car2	10	-	-	-	-	-	-	-	-	482.50
8	Dribble	2	-	0.23	0.42	-	-	-	1.75	0.90	1.08
12	Dribble	3	-	0.36	0.36	-	-	-	2.62	1.78	1.80
16	Dribble	4	-	0.51	0.51	-	-	-	6.94	3.23	3.14
20	Dribble	5	-	0.71	0.72	-	-	-	10.10	4.91	4.78
24	Dribble	6	-	0.92	0.93	-	-	-	16.85	7.16	7.04
28	Dribble	7	-	1.08	1.09	-	-	-	256.84	9.48	9.70
32	Dribble	8	-	1.64	1.73	-	-	-	84.88	12.90	13.39
36	Dribble	9	-	2.08	2.07	-	-	-	134.01	17.25	18.02
40	Dribble	10	-	2.74	2.76	-	-	-	135.79	23.07	23.49

TABLE II
RUNTIME (S) ON NONLINEAR INSTANCES. "-" INDICATES A TIMEOUT.

a nonlinear version of Dribble.

The Car domain includes only atomic actions and processes. The actions are to start or stop the Car, and accelerate or decelerate. The moving process models one-dimensional kinematics (distance as a function of velocity and velocity as a function of acceleration) and the wind-resistance process models the drag effect upon velocity. Additional actions for acceleration or deceleration increase the branching factor of the problem. As the problems scale, each instance i includes actions to accelerate and decelerate by $1, \ldots, i$ units. The linear and nonlinear versions of the domain differ in whether they include the nonlinear wind-resistance process.

The Generator domain includes two durative actions: generate, and refuel. The generate action has a duration of 1000 time units and consumes fuel at a linear rate. Its at-end effect satisfies the goal. Its overall condition requires that the fuel level is non-negative. The instances scale in the number of tanks required to refuel the Generator so that its overall condition is satisfied. The refuel actions increase the fuel level in the Generator continuously, by a linear rate (in the linear version) or a nonlinear rate (in the nonlinear version). For example, the refuel action defines the effects linearly as

```
(increase (fuel ?g) (* #t 2))
```

or nonlinearly

```
(increase (ptime ?t) (* #t 1))
(increase (fuel ?g)
          (* #t (* 0.1 (* (ptime ?t)
                          (ptime ?t)))))
```

The Dribble domain involves a process that effects the position x of a ball. The position changes continuously based upon the ball velocity v. The velocity changes continuously due to gravity ($-g$) and drag ($-0.1v^2$). The available actions are dribble(f) which decrease velocity by $f \in \{0, 1, 2, 4\}$. The dribble actions have the precondition that velocity is zero (i.e., the ball is at the top of its arc). The bounce event assigns velocity to $-0.9v$ and has the condition that the ball position x is zero. The initial state places the ball at $x = 1$ with velocity $v = 0$ and the goal is to reach $1.5 \leq x \leq 3.0$. The problem, while it does not scale, can be solved for plan lengths greater than one. We find plans (using two-step locking) for steps $k = 8, 12, 16, \ldots, 40$, which correspond respectively to $2, 3, 4, \ldots, 10$ dribble actions interleaved with the same number of bounce events.

Results: Tables I and II list runtime results for dReal and `HNSolve` on the respective linear and nonlinear instances. The columns list the number of encoding steps k, domain,

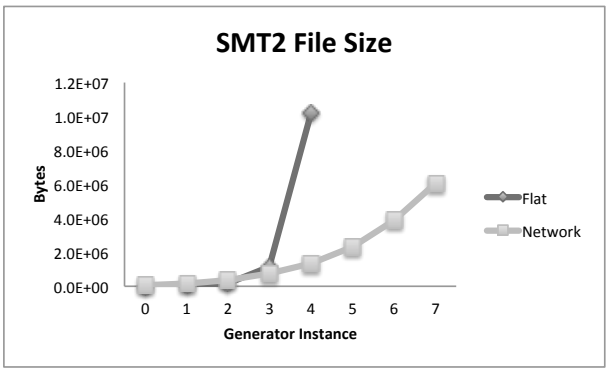

Fig. 1. Generator encoding file size.

instance, and run times in seconds for each encoding and solver configuration. The first three columns of results are denoted by "F" for a hand-coded flat encoding based upon the instances studied by Bryce *et al.* [10]. The second three columns of results are denoted by "C" for the automatically generated parallel composition of the network encoding into a single automaton. The last three columns denoted by "N" are the instances encoded with the network encoding. Within each group of columns, we denote by "F", "C", and "N" the results that do not use the new `HNSolve` layer. The columns with "+H" denote results for `HNSolve` without clause learning, and those with "+H+L", for `HNSolve` with clause learning. Entries with "-" indicate a timeout of 20 minutes was reached.

The results in the Generator domain are listed in the tables with the number of encoding steps k for the two-step lock encoding used in the C and N columns, and the steps in parentheses for the F encodings. The F encodings model the generate action with three steps, and each refuel action with two steps. The C and N encodings model each action with four steps, but is able to achieve the goal before releasing the final lock. Thus, each C and N instance uses one generate action (3 steps), and a number of refuel actions equal to the instance number (4 steps each). The linear results show that all solver configurations have difficulty scaling on the F and C encodings, as reported by [10]. The critical factor is that the encoding grows very large with the size of the instances (see Figure 1). Despite the poor scalability due to the size of the encoding, `HNSolve` (C+H, C+H+L) can provide some modest improvement over dReal (C). The network encoding N performs significantly better because it uses a tighter encoding. We also see the same trends as for the F and C encodings when comparing the different solver options; `HNSolve` (N+H, N+H+L) outperforms dReal (N) considerably and clause learning (+L) has a large impact.

The Car1 instances use the one-step lock and the Car2 instances use the two-step lock encoding. The F+H results on Car1 are most similar to the results reported by [10]. Our results are somewhat different because they are based upon dReal3, where the prior work results were collected on dReal2. The major difference between these versions of dReal are its use of the IBEX interval constraint solver (dReal3), and the Realpaver solver (dReal2). We see that the Car1 instances heavily favor the F encodings, and the use of dReal over `HNSolve`. In all caes, we see an improvement over dReal (F, C, N) by using `HNSolve` (+H). We see less of an improvement, and sometimes worse performance when using clause learning (+L). This may be due to the overhead associated with storing the clauses or a reshaping of the search space that leads to more backtracking.

The Car2 instances favor the N encoding and `HNSolve` with clause learning (N+H+L). It appears that the difference is that the relatively shorter encoding lengths in Car1 do not impact the encoding size as in Car2. In Car2 where the encoding length is double that of Car1, the `HNSolve` is needed to explore the search space. This result is largely consistent with the trend demonstrated in the Generator domain, where `HNSolve` performs best as the number of encoding steps increases.

Lastly, the Dribble domain highlights how both `HNSolve` and the network encoding have a positive impact upon performance. The F+H configuration and encoding is closest to that reported by Bryce *et al.* [10], and illustrates an improvement over previously published results that we attribute to a difference in dReal version.

Table III compares `HNSolve` (N+H+L) with the results reported by [10] (denoted as dReal2, and similar to F+H) and the other planners on linear instances of the Generator and Car domains. As above, each instance scales the respective number of tanks to fill the Generator (where each tank is required) and levels of acceleration/deceleration. We see that `HNSolve` scales much better than dReal on the Generator domain.

VI. Related Work

While PDDL+ [19] has been an accepted language for planning with continuous change for nearly a decade, very few planners have been able to handle its expressivity. Planners either assume that all continuous change is linear [7], [15], [16], [26] or handle nonlinear change by discretization [17].

LP-SAT [26] is very similar in spirit to our work because it uses a SAT solver to solve Boolean constraints and an LP solver to solve continuous (linear) constraints. The nature of the encodings is somewhat different in that our encoding makes use of the hybrid system semantics of PDDL+ in a very direct fashion. LP-SAT more closely resembles classical planning as SAT encodings. Unlike our work, LP-SAT does not incorporate heuristics.

More recent work [9], [13] has extended the LP-SAT approach by adapting its encoding for use in contemporary SMT solvers, including dReal and Z3 [24]. Unlike our work on `HNSolve`, these works focus solely on planning and not model checking hybrid systems. The advantage of focussing on planning encodings is that it is easier to implement the "must" semantics of PDDL+ and adapt existing techniques for SAT-based planning. Nevertheless, not all problems are best phrased as PDDL+, and approaches for reasoning about hybrid systems are necessary.

Dom	Planner	1	2	3	4	5	6	7	8
Gen	`HNSolve`	0.77	3.62	12.77	36.05	92.26	360.63	482.03	-
Gen	dReal2	3.07	15.6	134.71	1699.87	-	-	-	-
Gen	SpaceEx	0.01	0.03	0.07	0.1	0.19	0.28	0.45	0.65
Gen	CoLin	0.01	0.09	0.2	2.52	32.62	600.58	-	-
Gen	UPMur	0.2	18.2	402.34	-	-	-	-	-
Car	`HNSolve`	1.32	1.05	1.59	2.43	3.95	6.11	9.77	18.33
Car	dReal2	1.07	1.17	1.16	1.22	1.23	1.29	1.26	1.21
Car	SpaceEx	0.01	0.01	0.01	0.03	0.04	0.05	0.06	0.07
Car	CoLin	x	x	x	x	x	x	x	x
Car	UPMur	28.44	386.5	-	-	-	-	-	-

TABLE III
RUNTIME RESULTS (S) ON LINEAR GENERATOR AND CAR. "-" INDICATES A TIMEOUT.

Bogomolov *et al.* [7] and Della Penna *et al.* [17], like our work, make use of the planning as model checking paradigm. Unlike our work, Bogomolov *et al.* [7] encode a network of linear hybrid automata and we can handle nonlinear automata. Bogomolov *et al.* [7] use the SpaceEx model checker [20], which performs a symbolic search over the hybrid automata.

Coles *et al.* [16] and [15] approach PDDL+ from the perspective of heuristic state space search. Coles *et al.* [16] exploit piecewise linear representations of continuous change to derive powerful pruning conditions for forward heuristic search.

VII. CONCLUSION

We have described a new specialization of the dReal SMT solver called `HNSolve` and an associated network of hybrid automata encoding. The combination of `HNSolve` and the network encoding helps find PDDL+ plans as the number of encoding steps increases, especially with the use of a two-step lock encoding. We have shown that the approach scales up reasoning about PDDL+ planning and that it is competitive with the state of the art.

a) Acknowledgements: This work was supported under ONR contract N00014-13-1-0090.

REFERENCES

[1] R. Alur, T.A. Henzinger, G. Lafferriere, and G.J. Pappas, 'Discrete abstractions of hybrid systems', *Proceedings of the IEEE*, **88**(7), 971–984, (July 2000).

[2] Johan Bengtsson, Bengt Jonsson, Johan Lilius, and Wang Yi, 'Partial order reductions for timed systems', in *CONCUR'98 Concurrency Theory*, 485–500, Springer, (1998).

[3] S. Bogomolov, G. Frehse, R. Grosu, H. Ladan, A. Podelski, and M. Wehrle, 'A box-based distance between regions for guiding the reachability analysis of SpaceEx', in *Proceedings of the 24th International Conference on Computer Aided Verification (CAV 2012)*, eds., P. Madhusudan and Sanjit A. Seshia, volume 7358, pp. 479–494, (2012).

[4] Sergiy Bogomolov, Alexandre Donzé, Goran Frehse, Radu Grosu, Taylor T. Johnson, Hamed Ladan, Andreas Podelski, and Martin Wehrle, 'Guided search for hybrid systems based on coarse-grained space abstractions', *International Journal on Software Tools for Technology Transfer*, 1–19, (2015).

[5] Sergiy Bogomolov, Goran Frehse, Marius Greitschus, Radu Grosu, Corina S. Pasareanu, Andreas Podelski, and Thomas Strump, 'Assume-guarantee abstraction refinement meets hybrid systems', in *Hardware and Software: Verification and Testing - 10th International Haifa Verification Conference, HVC 2014, Haifa, Israel, November 18-20, 2014*, Lecture Notes in Computer Science, pp. 116–131. Springer, (2014).

[6] Sergiy Bogomolov, Daniele Magazzeni, Stefano Minopoli, and Martin Wehrle, 'PDDL+ planning with hybrid automata: Foundations of translating must behavior', in *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015, Jerusalem, Israel, June 7-11, 2015.*, eds., Ronen I. Brafman, Carmel Domshlak, Patrik Haslum, and Shlomo Zilberstein, pp. 42–46. AAAI Press, (2015).

[7] Sergiy Bogomolov, Daniele Magazzeni, Andreas Podelski, and Martin Wehrle, 'Planning as model checking in hybrid domains', in *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada.*, eds., Carla E. Brodley and Peter Stone, pp. 2228–2234. AAAI Press, (2014).

[8] Roberto Bruttomesso, Edgar Pek, Natasha Sharygina, and Aliaksei Tsitovich, 'The OpenSMT solver', in *Tools and Algorithms for the Construction and Analysis of Systems*, 150–153, Springer, (2010).

[9] Daniel Bryce, 'A happening-based encoding for nonlinear pddl+ planning', in *AAAI Workshop on Planning for Hybrid Systems*, (2016).

[10] Daniel Bryce, Sicun Gao, David J. Musliner, and Robert P. Goldman, 'SMT-based nonlinear PDDL+ planning', in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, eds., Blai Bonet and Sven Koenig, pp. 3247–3253. AAAI Press, (2015).

[11] Lei Bu, Alessandro Cimatti, Xuandong Li, Sergio Mover, and Stefano Tonetta, 'Model checking of hybrid systems using shallow synchronization', in *Formal Techniques for Distributed Systems*, ed., John Hatcliff; Elena Zucca, volume 6117 of *Lecture Notes in Computer Science*, 155–169, Springer, (2010).

[12] Lei Bu, You Li, Linzhang Wang, and Xuandong Li, 'BACH : Bounded reachability checker for linear hybrid automata', in *Formal Methods in Computer-Aided Design, 2008. FMCAD '08*, pp. 1–4, (Nov 2008).

[13] Michael Cashmore, Maria Fox, Derek Long, and Daniele Magazzeni, 'Full pddl+ planning through smt', in *AAAI Workshop on Planning for Hybrid Systems*, (2016).

[14] Alessandro Cimatti, Sergio Mover, and Stefano Tonetta, 'SMT-based verification of hybrid systems', in *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada.*, eds., Jörg Hoffmann and Bart Selman. AAAI Press, (2012).

[15] Amanda Coles, Andrew Coles, Maria Fox, and Derek Long, 'COLIN: Planning with continuous linear numeric change', *Journal of Artificial Intelligence Research*, **44**, 1–96, (2012).

[16] Amanda Jane Coles and Andrew Ian Coles, 'PDDL+ planning with events and linear processes', in *Proceedings of the Twenty-Fourth International Conference on Automated Planning and Scheduling, ICAPS 2014, Portsmouth, New Hampshire, USA, June 21-26, 2014*, eds., Steve Chien, Minh Binh Do, Alan Fern, and Wheeler Ruml. AAAI, (2014).

[17] Giuseppe Della Penna, Daniele Magazzeni, Fabio Mercorio, and Benedetto Intrigila, 'UPMurphi: A tool for universal planning on PDDL+ problems.', in *ICAPS*, pp. 106–113, (2009).

[18] Niklas Eén and Niklas Sörensson, 'An extensible SAT-solver', in *Theory and applications of satisfiability testing*, pp. 502–518. Springer, (2004).

[19] Maria Fox and Derek Long, 'Modelling mixed discrete-continuous domains for planning.', *J. Artif. Intell. Res.(JAIR)*, **27**, 235–297, (2006).

[20] G. Frehse, C. Le Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler, 'SpaceEx: Scalable verification of hybrid systems', in *Computer Aided Verification*, pp. 379–395, (2011).

[21] Sicun Gao, Jeremy Avigad, and Edmund M. Clarke, 'Delta-complete decision procedures for satisfiability over the reals', in *IJCAR*, pp. 286–300, (2012).

[22] Sicun Gao, Soonho Kong, and Edmund M. Clarke, 'dReal: An SMT solver for nonlinear theories over the reals', in *Automated Deduction - CADE-24 - 24th International Conference on Automated Deduction, Lake Placid, NY, USA, June 9-14, 2013. Proceedings*, ed., Maria Paola Bonacina, volume 7898 of *Lecture Notes in Computer Science*, pp. 208–214. Springer, (2013).

[23] Sicun Gao, Soonho Kong, and Edmund M Clarke, 'Satisfiability modulo odes', in *Formal Methods in Computer-Aided Design*, (2013).

[24] Leonardo Moura and Nikolaj Bjørner, *Tools and Algorithms for the Construction and Analysis of Systems: 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings*, chapter Z3: An Efficient SMT Solver, 337–340, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[25] Daniel Schwartz-Narbonne, Martin Schäf, Dejan Jovanovic, Philipp Rümmer, and Thomas Wies, 'Conflict-directed graph coverage', in *NASA Formal Methods - 7th International Symposium, NFM 2015, Pasadena, CA, USA, April 27-29, 2015, Proceedings*, eds., Klaus Havelund, Gerard J. Holzmann, and Rajeev Joshi, volume 9058 of *Lecture Notes in Computer Science*, pp. 327–342. Springer, (2015).

[26] Ji-Ae Shin and Ernest Davis, 'Processes and continuous change in a SAT-based planner', *Artif. Intell.*, **166**(1-2), 194–253, (2005).

[27] Pascal Van Hentenryck, David McAllester, and Deepak Kapur, 'Solving polynomial systems using a branch and prune approach', *SIAM Journal on Numerical Analysis*, **34**(2), 797–827, (1997).